

CIRCULATION COPY
SUBJECT TO RECALL
IN TWO WEEKS

UCRL- 91491
PREPRINT

DATABASE REQUIREMENTS FOR THE ADVANCED
TEST ACCELERATOR PROJECT

Frank W. Chambers

This paper was prepared for submittal to the
Hawaii International Conference on
System Sciences
Honolulu, Hawaii
January 2 - 4, 1985

November 5, 1984

Lawrence
Livermore
National
Laboratory

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

**DATABASE REQUIREMENTS FOR THE ADVANCED
TEST ACCELERATOR PROJECT***

**Frank W. Chambers
Lawrence Livermore National Laboratory
University of California
P.O. Box 808
Livermore, CA 94550**

November 5, 1984

ABSTRACT

The database requirements for the Advanced Test Accelerator (ATA) project are outlined. ATA is a state-of-the-art electron accelerator capable of producing energetic (50 million electron volt), high current (10,000 ampere), short pulse (70 billionths of a second) beams of electrons for a wide variety of applications. Databasing is required for two applications. First, the description of the configuration of facility itself requires an extended database. Second, experimental data gathered from the facility must be organized and managed to insure its full utilization. The two applications are intimately related since the acquisition and analysis of experimental data requires knowledge of the system configuration. This report reviews the needs of the ATA program and current implementation, intentions, and desires. These database applications have several unique aspects which are of interest and will be highlighted. The features desired in an ultimate database system are outlined.

*Work performed jointly under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract W-7405-ENG-48 and for the Department of Defense under Defense Advanced Research Projects Agency ARPA Order No. 4395 Amendment No. 31, Monitored by Naval Surface Weapons Center under document number N60921-84-WR-0080.

1. THE ADVANCED TEST ACCELERATOR - AN INTRODUCTION

1.1 The Advanced Test Accelerator Facility

The Advanced Test Accelerator (ATA)⁽¹⁾ is an electron accelerator constructed over the past four years at a cost of approximately \$50 million. The ATA was constructed by the Beam Research Group within M division of the Lawrence Livermore National Laboratory (LLNL). The primary sponsor for the project is the Defense Advanced Research Projects Administration (DARPA) with the Naval Surface Weapons Center (NSWC) acting as the DARPA agent.

The accelerator is remotely sited at Lawrence Livermore National Laboratory Site 300 about 15 miles southeast of the primary Livermore site. The accelerator is a linear machine and is 85 meters in length. A linear tunnel extends beyond the accelerator for another 115 meters. Currently, there is one alternate beam path available. Another such path into an experimental hall is under construction.

The accelerator is a linear induction machine of the type pioneered at Livermore and other places. This machine can produce a 10,000 ampere current pulse for 70 nanoseconds. The beam electrons are accelerated to 50 million volts or 100 times their rest mass; the beam is highly relativistic. The accelerator produces one pulse per second.

The ATA produced electron beam is unique due to its high current and high repetition rate. With the several beamlines and experimental chambers a variety of applications are being pursued. The propagation of the electron

beam through gas is being studied in detail. Radiation (both electromagnetic and ionizing) is produced when the beam strikes a beam dump. Radiation production can be measured and the radiation used in further experiments. The beam can be used as a driver for a free electron laser (FEL) experiment.⁽²⁾

1.2 ATA as an Information Source

The ATA requires information management in three main categories. First is the mechanical and electrical layout and interconnections of the machine. This configuration is constantly being altered as the accelerator is improved and the experiments changed. The time scale for changes is typically weeks to months. Coupled with the machine arrangement is the present values of numerous "softset" parameters such as power supply voltages and currents. These configuration values may change on a shot-to-shot basis. Finally, ATA is being used to conduct physics experiments. Much of the data is taken with digitizing hardware and stored in a digital form. The first two kinds of data constitute the machine configuration data and the third data category is referred to as the experimental data. At present about two megabytes of information are required to provide the hardware description of ATA. For the "softset" parameters about 32 kilobytes are required. Experimental data is produced at a rate of about 3 kilobytes/second for oscilloscope data and 500 kilobytes/second for television data. To make full utilization of the information generated by ATA requires proper data management. In this document we will discuss past experiences, present practices, and some of the needs and proposed implementations for a future system.

1.3 ATA Information Management

At present the ATA data management tasks are being performed in a piecemeal fashion. Components of the system include machine log books, experiment log books, digital data archives on several computers, a floppy disk library, and the collective memory of the experimentalists within the Beam Research group. Communication between various components of the system is performed manually by the person handling the data. Identifying and locating a data set is accomplished using log book records of data acquired. This is partially automated since a text editor can be used to view portions of the log book record. Machine configuration information is usually reconstructed after an experiment based on available records. Data is manipulated on one of several computers mentioned below. Every analysis and presentation technique discussed in this document is implemented on at least one computer system; some capabilities exist on all systems. This piecemeal approach has permitted much to be accomplished with limited resources. Much has been learned which will aid in the design of a fully integrated system which permit the identification, presentation, and analysis of all the ATA data within one framework.

Presently, the key element in the data management is the computer file. Experiments consist of collections of files. Files themselves are constructed of more fundamental building blocks. The sorting of experimental data takes on the file level. Files are identified for further analysis. This view of the data results in a very efficient division of labor. However, this approach does preclude the sorting of information based on the actual contents of a file and not merely several pointers. The full integration of our data

management system may proceed in two steps. First, a fully integrated system for accessing and sorting on all data files with the proper relations established between data and configuration must be approached. In this system the database manager would allow one to pinpoint files with information of interest. Analysis would then take place utilizing separate data analysis and manipulation codes. A second and potentially more interesting level of integration is possible, at least conceptually. A database manager could access the stored data at its basic structure. Sorting could occur on both acquisition attributes and data contents. Reports would consist not only of lists of files but also presentation of the contents of files.

A fully integrated system for data identification does not exist. At this stage in the design it is prudent to attempt to achieve a high degree of generality by viewing data in terms of its most primitive elements. As implementation nears a pruning of these ambitious goals will naturally occur.

1.4 ATA Computer Environment

A variety of computers are in use or available as indicated in Figure 1. Hardware is divided into several classes. Computers are located either at Site 300 or at the Livermore site. Computers may be dedicated to machine control, data acquisition, data analysis, data archiving, code development, or any combination of the above. The ATA accelerator is designed to be fully computer controlled. This system, running under a VAX-780, will both control and monitor the major mechanical and electrical components of the accelerator. System input is through touch panels and software selectable knobs. System

outputs control the vacuum, magnet, pulse power, and safety systems. In order to control the accelerator this system must have access to a full and current machine hardware description.

Data acquisition devices are either digitizing oscilloscopes or television systems. In the present configuration these devices are controlled directly by LSI-11 computers. A VAX-730 is available to operate in a supervisory capacity and as a user interface. Presently, the television systems run under the VAX, the oscilloscope control and acquisition system runs independent of the VAX. Signal selection, identification, display, recording, transcribing, and some limited analysis are carried out on this system.

The most powerful computer system available for Beam Research applications is the Octopus Network located at the Livermore site. Mainframes are a CRAY XMP, four CRAY 1s and 3 CDC 7600s. This system is available for time-sharing; time availability is limited but the large computational and storage capabilities of the system are sufficient to warrant its use for final analysis and archiving of the data. To maintain the security of information the Octopus network is a stand-alone system with hard-wired terminals. Modems, microwave links, and remote terminals are not permissible. Hence, data entry into Octopus must take place via magnetic tapes, floppy disks, or hand entry. This requirement does limit the utility of databases which cannot be shared between networks.

Finally, there are numerous personal computers within the beam research group. These small machines offer superior interaction to the time-shared system but limited numerical capabilities and storage capacity.

The location and accessibility of the configuration and experimental databases within this variety of computation hardware will be determined by several considerations. The on-site computers offer immediate response during the course of the experiment. However, their remote siting, limited access (machines are unavailable during data acquisition), and limited storage weigh against their use for the database. The Octopus system has virtually unlimited storage and extremely powerful numerical capabilities. However, interaction is sometimes slow and the system is not user friendly. Many experimental physicists will not use the Octopus system. Small computers offer excellent interactivity and impressive information handling capabilities. The cost and upper bound on storage are dropping at such a rate that the use of small machines in a future database management scheme must be considered seriously. Presently, the database management is taking place on all the above mentioned machines. The key element in developing analysis and reporting programs is the design of the information manipulation and user interface concepts. Once these concepts are crystalized through design and experience, implementation of a particular code on a particular machine is relatively straightforward. Data archiving is presently performed on the Octopus system. Data transmission is through floppy disks or to computers which are used as terminals.

1.5 Past History of Data Management

In designing the data management systems the customer talents must be taken into account. A historical review of past procedures is in order. Until the dawn of the 1980s most experimental data was taken using Polaroid

film as the recording medium. Data management consisted of gluing polaroids into a logbook. Data organization was strictly chronological. Experimental physicists become addicted to the smell of polaroid finishing chemicals. Machine configuration and signal identification were scratched into the film background. Experiments were carefully designed to highlight qualitative differences due to various effects. The need to digitize and perform detailed analysis of data was minimized. Experiments which required massive amounts of analysis of data were not performed.

The success of this historical approach cannot be overstated. Much was learned about beam propagation during this period. Data acquisition and analysis was cumbersome and limited; data management was not a problem.

The advent of digital data acquisition has proven a mixed blessing. One now has the ability to acquire massive amounts of data and to perform detailed and complex analysis of the data. Experiments which were previously too subtle and involved to be performed can now be addressed. In particular, full time resolution of beam quantities other than simple diagnostics can be accomplished. Data can be compared and correlated over longer time spans with greater ease. Data is more readily communicated to other theoretical and experimental physicists. However, now the entire process of data management is more complex. Automated data acquisition systems often obscure the results of the experiment from the physicist present when the data is taken. To ameliorate this problem a "quick look" capability has been incorporated into all data acquisition programs. The labeling of data is now far more difficult. The bulk of data is increased. The format of the data is largely determined by the acquisition systems and is not manifestly obvious from simply looking at data files. The experimentalist who labels his data by

scratching machine parameters onto his floppy disk cannot function in this environment. However, this same experimentalist with his experience and expertise must be fully integrated into the data management system.

2. MACHINE CONFIGURATION DATABASE

2.1 Applications and Contents

The machine description database serves three distinct purposes. First, mechanical and electrical engineers need to know at any given time what hardware is in the beam line and how it is connected. Second, the computer control system must have an image of the machine from which to perform its control and monitoring functions. Third, the diagnostic data acquisition system must have the capability of determining and recording the machine configuration at the time of the experiment for future analysis.

The machine configuration database will be used by the engineering staff to keep track of the machine configuration and parts inventory. Thus, the presence and physical location of beamline components must be maintained. Knowing component location will permit more efficient cabling of components. Having a list of all currently deployed beamline components will permit one to inventory remaining components more readily. Ideally, one would like to maintain a three-dimensional database on the machine configuration. Realistically, we intend to first implement a one-dimensional database with the dimension being distance along the accelerator. For the engineer the database should be up-to-date, the past history of the beamline is unimportant. From a hazard

control point of view, the past history of beamline components may be important as this determines their level of radioactivity.

The second area for a machine configuration database is the control system. Here the mechanical layout of the system is not as important as the electrical and plumbing connections to the system. Thus, to control the magnetic fields within the accelerator requires knowledge of the interconnections between the magnet power supplies and magnets. Vacuum system control depends on an adequate model of the different types of pumps and valves in use as well as the measurement gauges being employed. Control system information falls into two classes. There is hardset information about the equipment and connections such as which power supply drives which magnet. This information changes only on a several-month-time scale. Other parameters (such as the current being supplied to a particular magnet) constitute a softset of parameters. These values can and do change on a shot-to-shot basis.

The current version of the control system has a limited but functional machine image database. Due to time and personnel constraints, this machine description database may well remain independent of the one being constructed for engineering and experimental physics staff. This database must be up to date with the current configuration of the machine. Errors in this regard can be disastrous. The past history of the system wiring is not important.

The final application of the machine configuration database is in association with the experimental data acquisition. The proper classification of experimental data requires two types of information. The diagnostics placed along the beamline must be included in the machine configuration. The data acquisition system requires much information to access the proper signals.

Thus to attach a particular device to a particular input port, the system must know the proper wiring and/or switching required. To process the raw data of the measurement the system must have access to information regarding the details of the diagnostics (e.g., the cable lengths and probe calibrations). To determine the physical significance of the experimental data and to make cross comparisons with other data the entire machine hardware and software configuration must be available. Since data is continually being taken and archived the machine configuration at a particular time must be accessible. This presents one of the greatest challenges for the machine configuration database. Being up-to-date is not sufficient, past history must be easily accessible.

2.2 Machine Configuration Handling Present Practices

Machine configuration information is distilled from engineering drawings and notebooks, technicians reports, physicists' notes, and operational post mortems. Information is collected into "flat files" with fixed formats maintained using a simple editor. Everyday that the configuration changes a new file is created. Experimental data sets are manually associated with configuration files. These configuration files are neither complete nor accurate.

2.3 Features for a Future System

An extract from the machine configuration database shown in Figure 2. Components are uniquely identified by their beamline designation and location

within that beamline. Components are then characterized by various descriptive values. Characteristics pertaining to the installation, characteristics pertaining to the component class, and characteristics of this individual component within its class form separate tables within the database. Entries in the table would be ASCII strings, integer numbers, floating point numbers, and collections of numbers. The ability to deal with a collection of numbers (a TRACE) as an entity would be most helpful for characterizing magnets. A related useful entity would be a collection of numbers which represented vertices in line drawings of beamline components. This information fits well within the usual relational database. However, several unique features are desirable. Time is an important dimension within the database; one must be able to ascertain the machine state at any time in the past. Producing new copies of the database when parameters change is highly inefficient since parameters may change every second or every six months. The database might have an update capability where only update files are produced on the short-time scale. When a sufficient number of updates had occurred a new full database would be generated. This division of information should be transparent to the user. Inquiry into the database ideally should be on any computer or network available to the beam research personnel. Hence, the database structure should be as simple (even if rigid) as possible. Inquiry both by operators and other computer codes must be possible. Inquiry into past states of the machine must be feasible. Access to the database must be available at site 300 for the operations crew. Presently, this is accomplished with a room of blueprints and file cabinets. Such a system has worked with the usual problems concerning updating the information. Maintenance of the database should also occur at site. However, access to the database should be

available remotely. This access can be for read-only and need not be up to date within 24 hours.

Several types of inquiry are desirable. The layout of the beamline should be displayable from a starting location to an ending location. This display can be tabular or graphical. Reporting from the database with line drawings based on vertex information stored in component description tables is highly desirable. Tabulation of component type and location must be available. Nearest diagnostic should be accessible.

3. EXPERIMENTAL DATA DATABASE

3.1 Contents and Analysis

Data enters the system through several paths as indicated in Figure 3. At the signal level the configuration is highly experiment dependent. At the digitizer level the system is far more general. Data may be viewed as points, traces (oscilloscope data), arrays (TV pictures), and combinations of these entities. Point data can be probe positions, magnetic field settings, integrated charge on a beam current monitor etc. Such information enters the system from CAMAC crates, A/D converters, the control system, and the user. Trace information is presently acquired with Tektronix 7912 digitizing scopes. Typical sources are net current monitors (beambugs) which measure current flowing within a pipe in which the beam is propagating, radio frequency loops which detect the changing beam generated magnetic field, and photomultipliers which detect beam generated light or x-rays. The bandwidth of the oscilloscopes is 500 MHz so the system cannot be used for the signals with higher

frequency components. These high frequency signals will be acquired with television systems viewing the cathode ray tubes of fast oscilloscopes. Two-dimensional array information comes from television systems. Television systems are used to view light from any source including light generated when the beam strikes a foil, light from the beam passing through gas, and light from the accelerator cathode surface where the beam originates. Television information is either recorded on a video cassette recorder (VCR) or digitized on datacubes; spatial and temporal resolution are determined by the camera system in use. Framing cameras can be gated down to about 5 ns gate width. The streak camera will have a much higher time resolution.

Coupled with the data acquisition facilities at ATA are the data archiving, display and reduction capabilities which reside on the Octopus system with codes running on the CRAY computers. Data is archived permanently on the Octopus ATL (Automated Tape Library). Data resides at ATA for only a limited period. Data is transferred by disk or tape to the Livermore site; this transfer medium is saved as a backup to the Octopus system storage. Once the data from a major experiment is transferred to Octopus it is documented using the Octopus facilities for data reduction and output generation (i.e., data is put onto microfiche). Massive data reduction tasks for trace data are performed on this system. Here no data sorting takes place; all data is treated on an equal footing. The large amount of data available limits the usefulness of this type of reporting. At present we have image reduction codes running on CRAY computers but the actual, most efficient, mode of image data processing has yet to be determined. Octopus routines allow the full reduction of all the data in a timely and inexpensive fashion.

Data display and reduction also takes place on the data acquisition computers. This capability is an absolute necessity to aid the experimentalist in assessing the progress of an experiment. When the "quick look" facility is lacking the digital systems simply will not be used. Data reduction routines are now beginning to appear for the IBM Personal Computer. We find the massive data handling capabilities on the Octopus network and the interactive data handling capabilities on the smaller computers compliment each other quite well. It is highly desired to integrate these capabilities on one system.

3.2 Diagnostic Examples

An example of trace information is shown in Figure 4. Single traces are not usually taken; rather, sets of traces with some machine or diagnostic parameter varied are obtained. A typical example is beam current monitor data where the beam current and position are measured for all the beam current monitors along the accelerator. This collection of information allows one to determine the beam dynamics throughout the machine. These traces are then analyzed to obtain several key parameters to characterize the trace. For example, beam maximum current, beam charge (the time integrated current), beam arrival time, pulse rise length, and pulsewidth may be computed. One would like to characterize each trace with these "derived" parameters as well as the record time identification. However, the analysis to produce the "derived" values is often complex and ambiguous. The algorithms to compute these numbers are subtle and are often changed. The analysis is time consuming and often requires the use of very rapid and large computers. Thus one has a

database which contains information which one would like to be able to sort upon but which is not readily available.

Trace information may be obtained as machine or diagnostic conditions are discretely or continuously varied. The collection of beam current monitor data mentioned above is an example of discrete point collection. Trace data is taken as a function of probe position or magnetic field setting which can be continuously varied. Data is taken on a shot-by-shot basis. Now the interest is in examining the data versus probe position at a particular time. Traces are effectively assembled into a space-time picture of the beam. As seen in Figure 5, with a moving x-ray probe one obtains beam profile information at several times within the pulse. Now the "derived" parameters are the various radii of the beam as functions of time. Many single number parameterizations of the profile data are of value. Each experimental physicist has his favorite single number diagnostic for a profile; the database management system should be able to handle a significant number of such parameters.

Television data presents several problems for data management. First is the extreme bulk of data which must be recorded to preserve an image. With a 320 by 240 pixel array with a dynamic range of 256 (8 bits or 1 byte) a single frame requires 80 kilobytes of storage in absolute format and typically 250 kilobytes when put into ASCII. The acquisition of thousands of television pictures during a day is thus impossible. Secondly, the interpretation of images is very subtle; techniques painstakingly developed for one type of diagnostic are ineffectual for another type of image. The definition and computation of "derived" parameters is considerably more subtle than in the case of trace information.

The manipulation of TV data presents a challenge to the data acquisition system designers. Two ideas for improving the utility of TV data are being examined. The simple video cassette recorder (VCR) will be used as a secondary recording mechanism for TV data. To enhance the utility of the system digitizing frame grabbers will be used to hold the frame of interest for the one second between accelerator shots. Labels will be provided directly on the video signal recorded on the VCR. The VCR data will allow documentation of the experimental operations in a convenient and efficient manner. Subsequent digitization can be performed if desired; although with some loss of information. A database management system should have the capability of logging in and identifying the VCR tapes generated in the course of an experiment. Here the fundamental element to be identified would be the VCR tape and location on the tape. Anyone with a home recorder can appreciate the potential value of maintaining a database on VCR tape contents. A second option for improved utilization of TV data is to compress information during the acquisition phase and not attempt to store the entire image. If the "derived" parameters (such as beam radius, amplitude, baseline, and assymetry) could be extracted at acquisition time the bulk of data would be greatly reduced. At present users cannot agree on the proper algorithms for such information extraction and we are forced to record entire pictures.

3.3 Present Practices

At this point in time much thought is being given to future database capabilities. Present efforts are primitive.

Data sets are labeled with a signal source, a date, and a comment at acquisition time. Data sets are logged in when archived onto the Octopus system. A single data description file is maintained. Databasing performed either manually or with an editor which can perform pattern searches on this data description file. These implementations of databasing, while extremely primitive, have shown the value of data management for the experimental physicist. The need for more sophisticated and general systems has been demonstrated.

We are just beginning our search of existing packages for databasing. We have isolated several unique requirements or "wish list" desires which we now discuss.

3.4 Features for a Future Experimental Database

The ultimate form for the ATA experimental and configuration databases has not been determined. We can identify several capabilities which would be extremely useful in hardware and software. Present thinking has been that a database will handle pointers to individual files or blocks of data and that the numerical manipulation of this data will be done elsewhere. This division of information is occasioned by the limited storage and processing capabilities of currently available machines. In this section we will explore the features desirable in a database where storage and speed are assumed to be unlimited.

The database contains all the data information and not merely pointers to external data files. New entities, manipulations, and reporting modes are needed. A summary of the fundamental entities for our ultimate database are

shown in Figure 3. Our point of view changes from the database providing pointers to the relevant data blocks which are then analyzed with other codes to a fully integrated system where the database provides for the identification, location, extraction, manipulation, reporting, and correlating of data.

An additional fundamental entity for the database is the TRACE which consists of a collection of floating point numbers. TRACES will have parameters associated with them such as axis labels and sample intervals. Primitive manipulations can produce other traces or scalars. TRACE to TRACE manipulations include scalar multiplication, addition, subtraction, and fourier transforming. TRACE to scalar manipulations include minimum, maximum, span, and area. Several signal processing packages include many useful trace manipulations; each experimental physicist probably has his own package. The reporting of trace information will be by Plotting where the trace information is displayed graphically.

When the dimensionality of TRACE information is extended from 1 to 2 dimensions ARRAYS are created. ARRAYS of information can originate as pictures with a number of discrete pixels. ARRAYS can also be assembled as collections of TRACES. ARRAY manipulations include the usual arithmetic manipulations; for example, background subtraction.

The reporting of ARRAY information is by DISPLAYing as a picture, a contour plot, or a pseudo three-dimensional plot. TRACE information can be extracted from ARRAYS by slice extraction, integrated slice extraction, or various transforms. Single number diagnostics for ARRAYS also exist (e.g., span, volume, etc.).

Entities of higher dimensionality (e.g., MOVIES) can be conceived where the third dimension is spatial, temporal, or spectral (e.g., color). Even within the spirit of the presumptions of this section these entities are beyond current needs.

Databases often include fields which are functions of other fields within the data. This derived information can be computed when the data is originally entered (appended) into the database. However, due to the complexity of the algorithms for deriving information the database handling system will require a DIGESTION capability where missing derived elements within the database are calculated in a background mode. This is the converse of the usual "garbage collection"; the term "garbage generation" may be appropriate. The DIGESTION mode of operation could occur overnight when the database is not being queried. The user should be able to define the functions to be derived in terms of the primitive manipulations which can be performed on fields when constructing reports. Thus the database should have capabilities for field manipulation which resemble those of a spreadsheet. Simultaneously, since queries may occur when the data is not fully digested the system should provide for reporting when some information is missing.

Another fundamental concern in designing the integrated database is the historical dimension required for the machine configuration information. The history of the machine must be accessible. One approach is the creation of "daily" copies of the machine state. This is unappealing and wasteful of storage space. Another approach is the classification of parameters in terms of the time scales over which they may vary. The identification of hardset and softset parameters is a step in this direction. Past experience has shown that sooner or later the experiment will be done where a machine parameter

originally viewed as "cast in stone" will get changed on a shot-to-shot basis. Somehow a database is required which has updates and exceptions within the database framework.

The intent of the foregoing discussion has been that the database manager be sufficiently powerful to support all possible user inquiries. In this case the actual structure of the data within the database would be transparent to the user. In the real world it is preferable that the database structure be as simple as possible since alternate pathways into the database may be required. Not all queries into the database will occur through the database itself. Direct access to the data files in a simple fashion is desirable. For example, we have a very simple database implemented which characterizes beamline components. A second code permits the graphical display of beamline components. We would like the display code to be able to simply access the information maintained in our beamline database. Although we are working toward a fully integrated system; we require that some provision be made for access to the data outside the data manager.

4. MARKETABILITY

The databasing requirements of ATA are somewhat unique. One must ask if such a system is of broader interest. Anyone who owns a VCR can verify that sorting and location of information stored on VCRs is cumbersome. A database tailored to VCR usage would be of general interest in a wide market. The trace capabilities required in ATA are of interest to those who manipulate experimental data extensively. I do not see a broader market for this

capability. The ability to manipulate picture information has a wide variety of applications; particularly in the medical profession. A database which allows one to study diagnostic information (e.g., x-rays, sonograms) looking at time evolution, cross-correlations, and so forth would be most useful.

5. CONCLUSIONS

The maintenance and utilization of the ATA for physics experiments will be greatly facilitated by the development of machine configuration and experimental data databases. The requirements of these databases are unique and challenging. Their implementation is just beginning; future designs will depend critically on both operational experience and what is learned from the community.

REFERENCES

1. W. A. Barletta, "Energy and Technology Review," Lawrence Livermore National Laboratory, UCRL-52000-79-9, September 1979.
2. D. Prosnitz, "Energy and Technology Review," Lawrence Livermore National Laboratory, 6CRL-52000-82-1, January 1982.
3. F. W. Chambers, J. Kallman, J. McDonald and M. Slominski, "ATA Diagnostic Data Handling System—An Overview," Lawrence Livermore National Laboratory, UCRL-90912, June 1984.

ACKNOWLEDGEMENT

Jeff Kallman and Jim McDonald have contributed significantly both in developing the existing data acquisition systems and exploring the concepts for the next generation system. All the users of the existing systems have contributed ideas, requests, and unreasonable demands for the next system.

FWC:emj

8625v/0160v

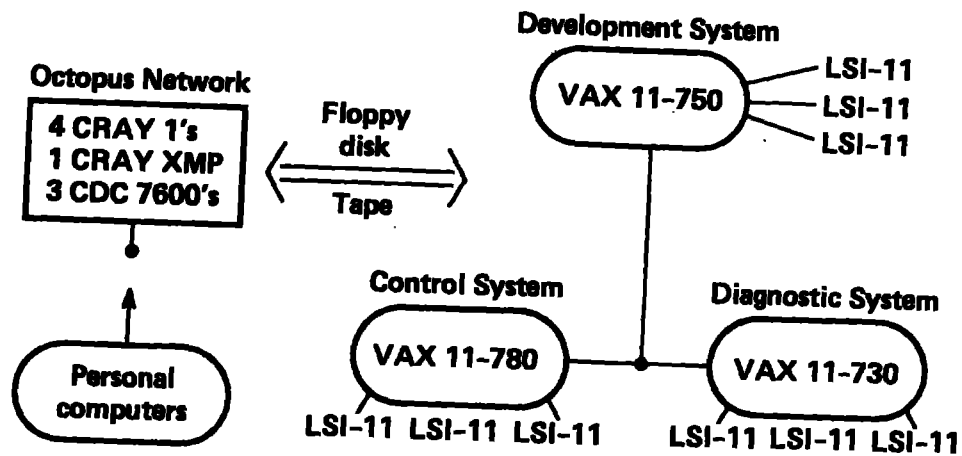


Fig. 1. Computer hardware available to the Beam Research Group for ATA machine control, data acquisition, and data analysis. Communications between systems are indicated.

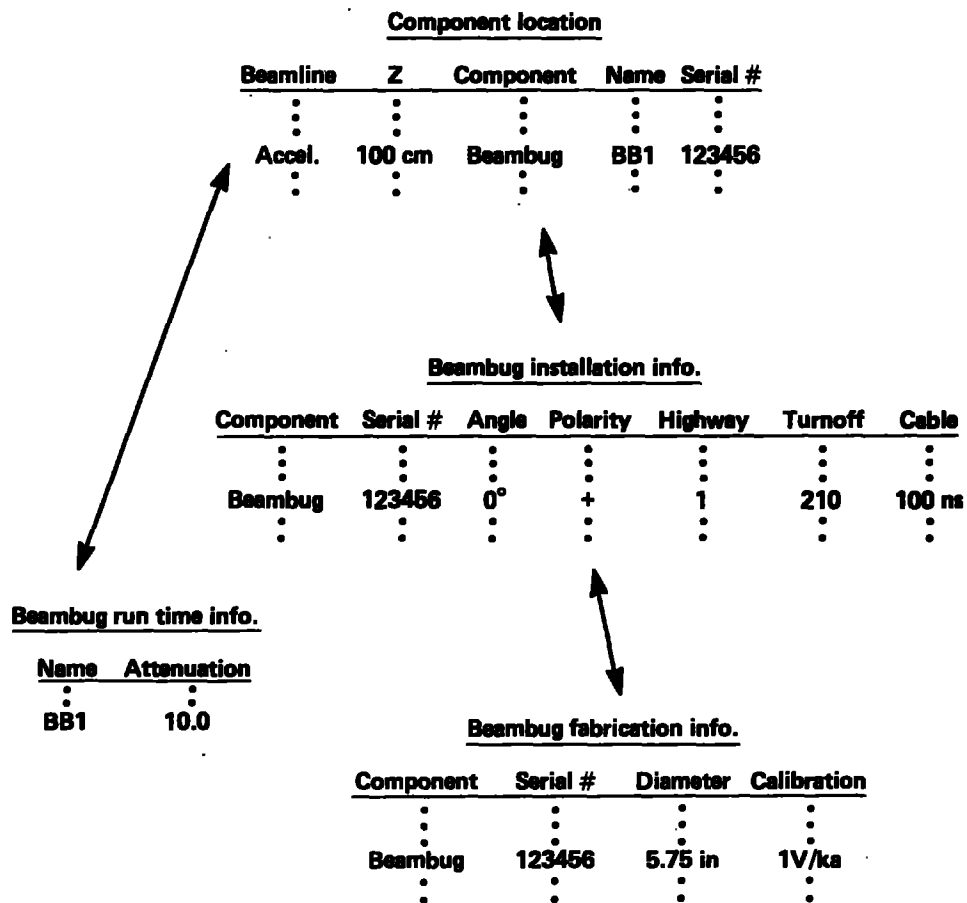


Fig. 2. A portion of the proposed configuration database structure indicating the relations among various components. Information on a beam current monitor (beambug) falls into the categories of fabrication, installation, location, and utilization.

| Signal | Digitizer | Data form |
|-----------------------|---|----------------------|
| Probe position | CAMAC Control system User | f point |
| Field current | | |
| Cable length | | |
| • | | |
| • | | |
| • | | |
| Beambugs I, X, Y | 7912 oscilloscope 7104 oscilloscope + TV | f _i trace |
| RF loops T, B, E, W | | |
| Photomultiplier tubes | | |
| • | | |
| • | | |
| Foil light | TV camera or streak camera + datacube | f _i array |
| Gas light | | |
| Cathode light | | |
| Surveillance | | |
| • | | |
| • | | |

Data is best viewed as a collection of fundamental entities:




Fig. 3. Data inputs to the diagnostic computer system. A wide variety of signals lead to data which fall into four very simple classes.




Fig. 4. An example of trace information as displayed on an IBM Personal Computer. All comments are included on this report as well as the derived parameters; Ymin, Ymax, and Area. At this point no analysis has been performed.

DATAOUT : P19485 BOW PROBE .28 M .08T AIR/50T NEO 23-MAY-83
TIMES : 1.50E+01 1.56E+01 1.62E+01 1.69E+01

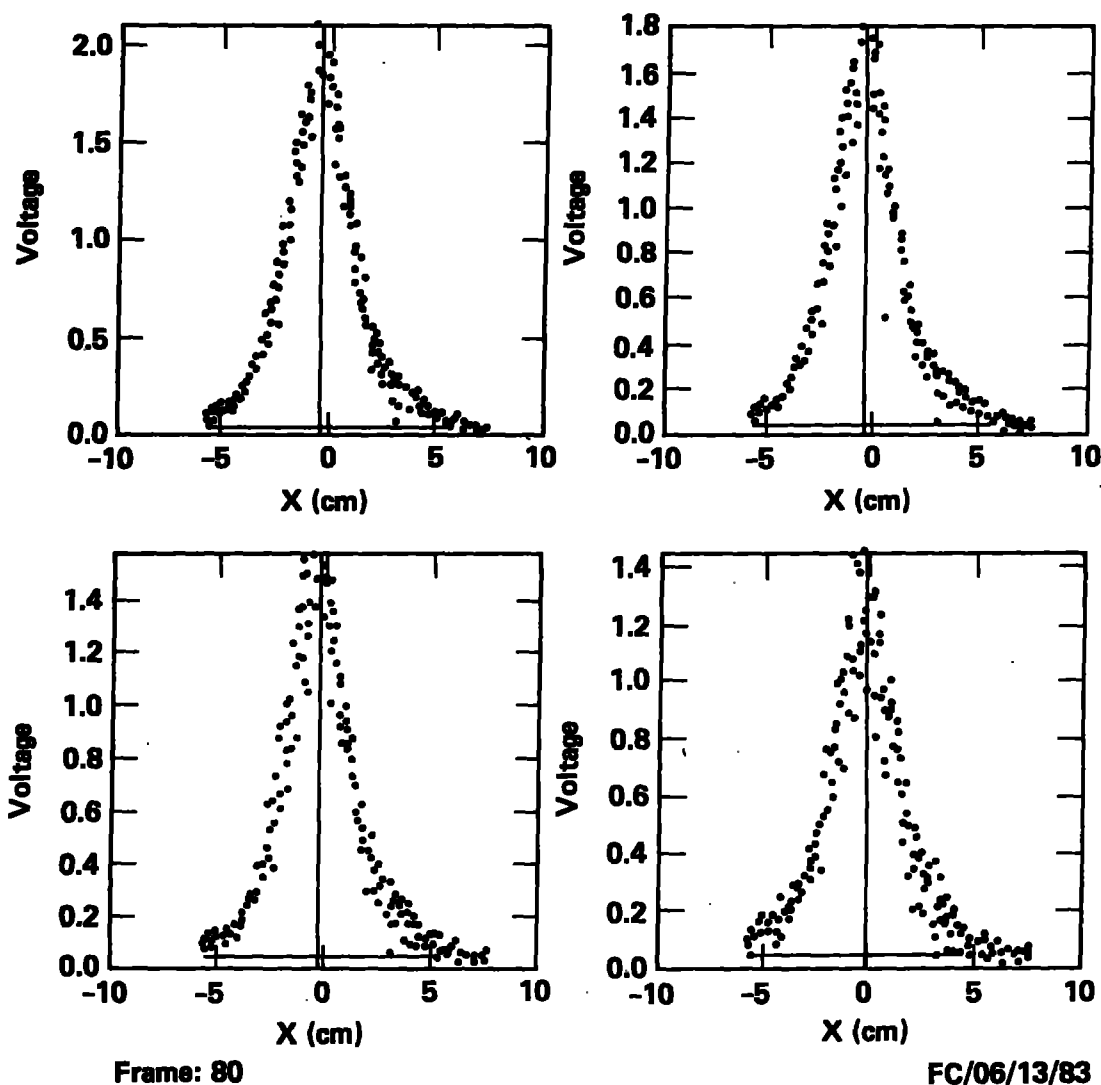


Fig. 5. An example of probe data indicating the beam profile at several times. This display was assembled on a CRAY computer. Lines on the data plot indicate the derived parameters of profile baseline and center.